

Pareto-Path Multi-Task Multiple Kernel Learning

Cong Li, Michael Georgiopoulos and Georgios C. Anagnostopoulos

congli@eecs.ucf.edu, michaelg@ucf.edu and georgio@fit.edu

Keywords: Multiple Kernel Learning, Multi-task Learning, Multi-objective Optimization, Pareto Front, Support Vector Machines

Abstract

A traditional and intuitively appealing Multi-Task Multiple Kernel Learning (MT-MKL) method is to optimize the sum (thus, the average) of objective functions with (partially) shared kernel function, which allows information sharing amongst tasks. We point out that the obtained solution corresponds to a single point on the Pareto Front (PF) of a Multi-Objective Optimization (MOO) problem, which considers the concurrent optimization of all task objectives involved in the Multi-Task Learning (MTL) problem. Motivated by this last observation and arguing that the former approach is heuristic, we propose a novel Support Vector Machine (SVM) MT-MKL framework, that considers an implicitly-defined set of conic combinations of task objectives. We show that solving our framework produces solutions along a path on the aforementioned PF and that it subsumes the optimization of the average of objective functions as a special case. Using algorithms we derived, we demonstrate through a series of experimental results that the framework is capable of achieving better classification performance, when compared to other similar MTL approaches.

1 Introduction

Multiple Kernel Learning (MKL) is an important method in kernel learning that has drawn considerable attention since being introduced in [19]. MKL seeks an appropriate kernel function (and, hence, kernel matrix) by linearly or non-linearly combining several pre-selected candidate kernel functions. Given a machine learning problem/task, the optimal kernel matrix is derived by optimizing the associated objective function with respect to the combination coefficients, say $\theta = [\theta_1, \dots, \theta_M]$, for M pre-specified kernels. Employing MKL avoids having to resort to (cross-)validating over different kernel choices, which is computationally cumbersome and may require additional data for validation. A key focus of MKL is identifying solutions for θ by imposing appropriate constraints on it, such as an L_1 -norm constraint [19] [23], L_2 -norm constraint [17], and L_p -norm constraint with $p > 1$ [18] as a generalization of the previous two methods. The generalization bound and other theoretical aspects of the L_p -norm MKL method is extensively studied in [16]. Besides searching for the optimal constraints on θ , several other works have been proposed, such as using a Group-Lasso type regularizer [35] and an L_1 -norm within-group / L_s -norm ($s \geq 1$) group-wise regularizer [1], nonlinearly combined MKL [6], MKL with localized θ [10], MKL with hyperkernels [22], MKL based on the radii of minimum enclosing balls [9] and other methods, such as the ones of [23], [29] and [33], to name a few. A thorough survey of MKL is given in [11].

Another active path of MKL research is combining MKL with Multi-Task Learning (MTL), which is commonly referred to as Multi-Task Multiple Kernel Learning (MT-MKL). MTL aims to simultaneously learn multiple related tasks using shared information, such that each task can benefit from learning all tasks. Existing approaches consider several different types of information sharing strategies. For example, [1], [14] and [24] applied a mixed-norm regularizer on the weights of each linear model (task), which forces tasks to be related, and, at the same time, achieves different levels of inner-task and inter-task sparsity on the weights. Another example is the model proposed in [34], which considers T tasks and restricts the T Support Vector Machine (SVM) weights to be close to a common weight, such that the weights from all tasks are

related. Additionally, for the recently proposed *Minimax MTL* model [20], tasks are related by minimizing the maximum of the T loss functions, in order to guarantee some minimum level of accuracy for each task. Last but not least, a straightforward strategy of information sharing is to let all tasks share (or partially share) a common kernel function, which has been investigated in [15], [30], and [25], again to name a few. According to this strategy, tasks are related by mapping data from all tasks to a common feature space and, subsequently, each task is learned in the same, common feature space.

Considering the latter strategy of information sharing, the most intuitively straightforward formulation is to optimize the sum (equivalently, the average) of objective functions with shared kernel function, such as the models in [30] and [25]. However, as we subsequently argue, this method is rather an *ad hoc* strategy. First, we observe that optimizing the average of objective functions is equivalent to finding a particular solution to a Multi-Objective Optimization (MOO) problem, which aims to optimize all task objectives simultaneously. In specific, it is a well known fact (see [4, p. 178]) that scalarizing a MOO problem by optimizing different conic combinations (linear combinations with non-negative coefficients) of the objective functions leads to the discovery of solutions that correspond to points on the convex (when minimizing) part of the problem's Pareto Front (PF). The latter set is the set of non-dominated solutions in the space of objective values. Hence, by optimizing the average of task objectives in an MTL setting, one only finds a particular PF point (or more, if the PF is non-convex) of the corresponding MOO problem. Therefore, while considering the optimization of this average may be intuitively appealing, it is, nevertheless, a largely *ad hoc* strategy. Foreseeably so, optimizing a different conic combination of task objectives, albeit among an infinity of possibilities, may improve the performance of each task even further, when compared to the case of averaged objectives. This amounts to searching for better solutions, *i.e.* PF points, of the associated MOO problem and forms the basis of our work.

In this paper, we propose a new SVM-based MT-MKL framework for binary classification tasks. The common kernel utilized by these SVM models is established through a typical MKL approach. More importantly, though, it considers optimizing specific conic combinations of the task objectives, in order to improve upon the traditional MTL method of averaging. In Section 2 we show that the obtained solutions trace a path on the PF of the relevant MOO problem. While doing so does not explore the entire PF, searching for solutions to our problem is computationally feasible. The framework's conic combinations and, thus, the aforementioned path, is parameterized by a parameter $p > 0$. For $p = 1$, the whole framework coincides with the traditional MTL approach of minimizing the average of SVM objective functions. In Section 3 and Section 4 we derive algorithms to solve our proposed framework for $p \geq 1$ and $0 < p < 1$ respectively, while in Section 5 we demonstrate the impact of p 's value on learning performance. In specific, we show that recognition accuracy for all tasks increases as p decreases below 1. We discuss why this phenomenon occurs and provide insights into the behavior of our MT-MKL formulation. In the same section, we also provide a variety of experimental results to highlight the utility of this formulation. Finally, we briefly summarize all our findings in Section 6.

In the sequel, we'll be using the following notational conventions: vector and matrices are denoted in boldface. Vectors are assumed to be columns vectors. If \mathbf{v} is a vector, then \mathbf{v}' denotes the transposition of \mathbf{v} . Vectors $\mathbf{0}$ and $\mathbf{1}$ are the all-zero and all-one vectors respectively. Also, \succeq and $\max\{\cdot, \cdot\}$ between vectors will stand for the component-wise \geq and $\max\{\cdot, \cdot\}$ relations respectively. For any $\mathbf{v} \succeq \mathbf{0}$, \mathbf{v}^p represents component-wise exponentiation of \mathbf{v} . Furthermore, we will be using the notation $\nu(\mathbf{v})_p \triangleq (\mathbf{1}'\mathbf{v}^p)^{\frac{1}{p}}$ where \mathbf{v} is a vector with $\mathbf{v} \succeq \mathbf{0}$ and $p \in (0, +\infty]$. Observe that for $p \geq 1$, $\nu(\mathbf{v})_p = \|\mathbf{v}\|_p$, where $\|\cdot\|_p$ stands for the ordinary Minkowski L_p -norm for finite-dimensional vectors. Note that for $p \in (0, 1)$, $\nu(\mathbf{v})_p$ is not a norm. For any $s > 0$ and vector \mathbf{v} , we define the set $\bar{B}_{\mathbf{v},s} \triangleq \{\mathbf{v} | \mathbf{v} \succeq \mathbf{0}, \nu(\mathbf{v})_s \leq 1\}$. Also, let Z_n be the set of integers $\{1, \dots, n\}$ for any $n \geq 1$. Finally, notice that the proofs of the manuscript's theoretical results are provided in the Appendix.

2 Problem Formulation

Consider the following MT-MKL problem involving T SVM training tasks:

$$\begin{aligned}
& \min_{\mathbf{f}, \boldsymbol{\theta}, \boldsymbol{\xi}, \mathbf{b}} \nu(\mathbf{g}(\mathbf{f}, \boldsymbol{\theta}, \boldsymbol{\xi}))_p \\
& \text{s.t. } y_i^t \left(\sum_{m=1}^M f_m^t(\mathbf{x}_i^t) + b^t \right) \geq 1 - \xi_i^t, \quad \forall i \in Z_{N_t}, t \in Z_T \\
& \quad \boldsymbol{\xi}^t \succeq \mathbf{0}, \quad \forall t \in Z_T \\
& \quad \boldsymbol{\theta} \in \bar{B}_{\boldsymbol{\theta}, s}, \quad s \geq 1
\end{aligned} \tag{1}$$

where $\mathbf{g}(\mathbf{f}, \boldsymbol{\theta}, \boldsymbol{\xi}) \triangleq [g^1(\mathbf{f}, \boldsymbol{\theta}, \boldsymbol{\xi}), \dots, g^T(\mathbf{f}, \boldsymbol{\theta}, \boldsymbol{\xi})]'$ and each $g^t(\mathbf{f}, \boldsymbol{\theta}, \boldsymbol{\xi})$ is defined as the t -th multi-kernel SVM objective, *i.e.*

$$g^t(\mathbf{f}, \boldsymbol{\theta}, \boldsymbol{\xi}) \triangleq \sum_{m=1}^M \frac{\|f_m^t\|_{H_m}^2}{2\theta_m} + C \sum_{i=1}^{N_t} \xi_i^t \tag{2}$$

where $\mathbf{f} \triangleq [\mathbf{f}^1, \dots, \mathbf{f}^T]'$, $\mathbf{f}^t \triangleq [f_1^t, \dots, f_M^t]'$, $\boldsymbol{\xi} \triangleq [\boldsymbol{\xi}^1, \dots, \boldsymbol{\xi}^T]'$, $\boldsymbol{\xi}^t \triangleq [\xi_1^t, \dots, \xi_{N_t}^t]'$, $\boldsymbol{\theta} \triangleq [\theta_1, \dots, \theta_M]'$, $\mathbf{b} \triangleq [b^1, \dots, b^T]'$. Moreover, $\{\mathbf{x}_i^t, y_i^t\}$, where $i = 1, \dots, N^t$, are the training samples available for the t^{th} task.

For each task t , M discriminative functions f_m^t are sought under the constraints $f_m^t \in \mathcal{H}_m$, where each \mathcal{H}_m is a Reproducing Kernel Hilbert Space (RKHS) associated to a pre-selected reproducing kernel $k_m(\cdot, \cdot)$. Let \mathbf{x} encompass all the optimization variables, that is, $\mathbf{f}, \boldsymbol{\theta}, \boldsymbol{\xi}, \mathbf{b}$. We can restate Problem (1) as

$$\min_{\mathbf{x} \in \Omega(\mathbf{x})} \nu(\mathbf{g}(\mathbf{x}))_p \tag{3}$$

where $\Omega(\mathbf{x})$ is the feasible region of \mathbf{x} , given by the constraints of Problem (1). Note that the SVM objective functions $\mathbf{g}(\mathbf{x})$ are non-negative and not all simultaneously zero for any value of \mathbf{x} . Therefore, $\nu(\mathbf{g}(\mathbf{x}))_p$ is well defined, based on the definition of $\nu(\cdot)_p$ given in Section 1.

Since $\nu(\cdot)_1 = \|\cdot\|_1$, the traditional MT-MKL formulation considers only the case, when $p = 1$, *i.e.*, the sum (or, equivalently, the average) of the T objectives. However, as argued in Section 1, optimizing the objectives' average may in practice not necessarily lead to achieving the best obtainable performance for every task simultaneously and, thus, it is of interest to investigate cases, for which $p \neq 1$. We show that, for any $p > 0$, the optimum value \mathbf{g}^* is a PF solution for the T SVM objectives. Based on this conclusion, we are able to explore a path on the PF by tuning only one parameter, namely p . We later discuss that by doing so, it not only helps us achieve uniform performance improvements, but it also provides useful insights into the SVM-based MT-MKL problem we are considering.

Proposition 1. *For any $p > 0$ and arbitrary vector function $\mathbf{g}(\mathbf{x}) \in \mathbb{R}^T$ with $\mathbf{g} \succeq \mathbf{0}$ for all vectors \mathbf{x} in a feasible set $\Omega(\mathbf{x})$, the optimal solution \mathbf{x}^* of the general optimization problem*

$$\min_{\mathbf{x} \in \Omega(\mathbf{x})} \nu(\mathbf{g}(\mathbf{x}))_p \tag{4}$$

is a PF solution of the MOO problem

$$\min_{\mathbf{x} \in \Omega(\mathbf{x})} \mathbf{g}(\mathbf{x}) \tag{5}$$

which considers the simultaneous minimization of the T objectives $g_1(\mathbf{x}), \dots, g_T(\mathbf{x})$.

The proof is given in Section .1 of the Appendix. It is readily evident that our framework is convex, when $p \geq 1$, and non-convex, when $p < 1$. In the following two sections, we discuss how to optimize Problem (1) in both cases.

3 Learning in the Convex Case

For $p \geq 1$, we first convert Problem (1) to an equivalent min-max optimization problem and then employ a specialized version of an algorithm proposed in [32] to solve it.

Lemma 1. Let $p \geq 1$, $\lambda, g \in \mathbb{R}^T$ such that $g \succeq \mathbf{0}$, but $g \neq \mathbf{0}$. Also, let $q \triangleq \frac{p}{p-1}$. Then,

$$\max_{\lambda \in \bar{B}_{\lambda, q}} \lambda' g = \nu(g)_p = \|g\|_p \quad (6)$$

Furthermore, a solution to the previously stated maximization problem is given as

$$\lambda^* = \begin{cases} (\frac{g}{\|g\|_p})^{p-1} & \text{if } p > 1 \\ \mathbf{1} & \text{if } p = 1 \end{cases} \quad (7)$$

The veracity of the previous lemma is easily demonstrated. In specific, note that Equation (6) is similar to the definition of the dual norm (see [4, p. 637]):

$$\|g\|_* = \sup_{\|\lambda\| \leq 1} \lambda' g \quad (8)$$

where $\|\cdot\|_*$ is the dual norm of $\|\cdot\|$. A slight difference between Problem (6) and Problem (8) is that, in Problem (6), the constraint is $\lambda \in \bar{B}_{\lambda, q}$ (i.e. $\lambda \succeq \mathbf{0}$ and $\|\lambda\|_q \leq 1$). However, note that, as $g \succeq \mathbf{0}$, the optimal λ must satisfy $\lambda \succeq \mathbf{0}$. Therefore, under the condition $g \succeq \mathbf{0}$, (6) is the same as the definition of the dual norm. Finally, Equation (7) gives the solution of the maximization problem in Equation (6) and its correctness can be verified by directly substituting Equation (7) into Problem (6).

The previous lemma implies that, for $p \geq 1$, instead of optimizing $\nu(g)_p$, one could optimize the conic combination of objective functions g_t with coefficients λ^* as given in Equation (7). Furthermore, it allows us to transform Problem (1) into an equivalent, more suitable problem via the next theorem.

Theorem 1. For $p \geq 1$, MT-MKL optimization Problem (1) is equivalent to the following min-max problem

$$\begin{aligned} \min_{\theta} \max_{\lambda, \beta} \quad & \Phi(\theta, \beta, \lambda) \\ \text{s.t.} \quad & \beta^{t'} \mathbf{y}^t = 0, \quad \lambda^t C \mathbf{1} \succeq \beta^t \succeq \mathbf{0}, \quad t \in Z_T; \\ & \theta \in \bar{B}_{\theta, s}, \quad s \geq 1; \quad \lambda \in \bar{B}_{\lambda, q}. \end{aligned} \quad (9)$$

where $\Phi(\theta, \beta, \lambda) \triangleq \sum_{t=1}^T (\beta^{t'} \mathbf{1} - \frac{1}{2\lambda^t} \beta^{t'} \mathbf{Y}^{t'} (\sum_{m=1}^M \theta_m \mathbf{K}_m^t) \mathbf{Y}^t \beta^t)$ and $q \triangleq \frac{p}{p-1}$.

The proof of the above theorem is given in Section 2 of the Appendix. In Problem (9), $\beta_i^t \triangleq \alpha_i^t \lambda^t$, where α_i^t 's are the dual variables of the t -th SVM problem, $\mathbf{y}^t \in \{-1, 1\}^{N_t}$ is the vector containing all labels of the training data for the t -th task, \mathbf{Y}^t is a diagonal matrix with the elements of \mathbf{y}^t on its diagonal, and \mathbf{K}_m^t is the kernel matrix with elements $k_m(\mathbf{x}_i^t, \mathbf{x}_j^t)$.

3.1 Tseng's Algorithm

Note that Problem (9) is a convex-concave optimization problem. Based on this fact, we consider Tseng's algorithm [32] for solving the problem. Define $\mathbf{u} \triangleq [\theta', \beta', \lambda']'$ and let $\Phi(\mathbf{u}) \triangleq \Phi(\theta, \beta, \lambda)$ stand for the objective function of Problem (9). Moreover, the algorithm considers the vector function $\mathbf{q}(\mathbf{u}) \triangleq [\nabla_{\theta} \Phi(\mathbf{u})', -\nabla_{\beta} \Phi(\mathbf{u})', -\nabla_{\lambda} \Phi(\mathbf{u})']'$. During the k -th iteration, assuming that \mathbf{u}_k is already known, the algorithm finds $\zeta > 0$, such that \mathbf{v}_k , which is given by

$$\mathbf{v}_k = \arg \min_{\mathbf{u} \in \Omega(\mathbf{u})} \{\mathbf{u}' \mathbf{q}(\mathbf{u}_k) + \zeta D(\mathbf{u}, \mathbf{u}_k)\} \quad (10)$$

satisfies the following condition:

$$\min_{\mathbf{u} \in \Omega(\mathbf{u})} \{\mathbf{u}' \mathbf{q}(\mathbf{v}_k) + \zeta D(\mathbf{u}, \mathbf{u}_k)\} \geq \mathbf{v}_k' \mathbf{q}(\mathbf{v}_k) \quad (11)$$

Subsequently, \mathbf{u}_{k+1} is set as the minimizer of problem (11). In the last two problems, $D(\cdot, \cdot)$ denotes the Bregman divergence, which, for any strictly convex function h , is defined as $D(\mathbf{u}, \mathbf{v}) = h(\mathbf{u}) - h(\mathbf{v}) - (\mathbf{u} - \mathbf{v})' \nabla h(\mathbf{v})$. $\Omega(\mathbf{u})$ is the feasible region of \mathbf{u} ; for our problem it is the feasible set of θ, β and λ given by the constraints in Problem (9). To find ζ , the authors in [21] suggest initializing ζ to a large positive value and

then to keep halving it until (11) is satisfied. Tseng's algorithm successfully converges, when Φ is convex-concave and differentiable with Lipschitz-continuous gradient \mathbf{q} . This condition is not satisfied, when $\exists t$ such that $\lambda_t = 0$. However, we will show that this will never happen to our algorithm and, thus, it does not affect its convergence.

3.2 Adaptation of Tseng's Algorithm to our Framework

In order to solve Problem (9) with Tseng's algorithm, which consists of solving Problem (10) and the minimization problem on the left side of Problem (11), we first show how to solve Problem (10) in our setting. We choose $h(\mathbf{u}) = h_{\theta}(\theta_{\mathbf{u}}) + h_{\beta}(\beta_{\mathbf{u}}) + h_{\lambda}(\lambda_{\mathbf{u}})$, where $h_{\beta}(\beta) \triangleq \|\beta\|_2^2$, $h_{\lambda}(\lambda) \triangleq \|\lambda\|_2^2$, and $h_{\theta}(\theta) \triangleq \frac{1}{\bar{s}} \|\theta\|_{\bar{s}}^{\bar{s}}$ with $\bar{s} = s$, when $s > 1$, and $\bar{s} = 2$, when $s = 1$. Given our choices, it is not difficult to see that the minimizations in Problem (10) can be separated into the following two problems:

$$\min_{\lambda, \beta \in \Omega(\lambda, \beta)} \|\beta\|_2^2 - \beta' \left(\frac{1}{\zeta} \nabla_{\beta} \Phi(\mathbf{u}_k) + 2\beta_{\mathbf{u}_k} \right) + \|\lambda\|_2^2 - \lambda' \left(\frac{1}{\zeta} \nabla_{\lambda} \Phi(\mathbf{u}_k) + 2\lambda_{\mathbf{u}_k} \right) \quad (12)$$

$$\min_{\theta \in \Omega(\theta)} \frac{1}{\bar{s}} \|\theta\|_{\bar{s}}^{\bar{s}} - \theta' \left(\theta_{\mathbf{u}_k}^{\bar{s}-1} - \frac{1}{\zeta} \nabla_{\theta} \Phi(\mathbf{u}_k) \right) \quad (13)$$

Problem (12) is a convex optimization problem, which can be solved via many general-purpose optimization tools, such as `cvx` [13][12]. On the other hand, Problem (13) has a closed-form solution that is provided by the following theorem.

Theorem 2. *Let $\theta \in \mathbb{R}^M$ and $\psi \in \mathbb{R}^M$ such that $\psi \succeq \mathbf{0}$, $s > 1$ and $r \triangleq \frac{1}{s-1}$. The unique solution of the constrained minimization problem*

$$\min_{\theta \in \Omega(\theta)} \frac{1}{s} \mathbf{1}' \theta^s - \psi' \theta \quad (14)$$

is given as

$$\theta^* = \frac{\psi^r}{\max \{1, \|\psi^r\|_s\}} \quad (15)$$

if $\Omega(\theta) \triangleq \{\theta | \theta \in \bar{B}_{\theta, s}\}$ and

$$\theta^* = (\max \{\psi - \mu \mathbf{1}, \mathbf{0}\})^r \quad (16)$$

if $\Omega(\theta) \triangleq \{\theta | \theta \in \bar{B}_{\theta, 1}\}$. In Equation (16), μ is the smallest nonnegative real number, such that $\|\theta^*\|_1 = 1$.

Solving the minimization problem in (11) can be accomplished by utilizing the similar procedure as the one for solving Problem (10), since these two problems have the same form. The algorithm is deemed to have converged, when the duality gap $\max_{\beta, \lambda \in \Omega(\beta, \lambda)} \Phi(\theta_{\mathbf{u}_k}, \beta, \lambda) - \min_{\theta \in \Omega(\theta)} \Phi(\theta, \beta_{\mathbf{u}_k}, \lambda_{\mathbf{u}_k})$ is smaller than a predefined threshold. The algorithm is summarized in Algorithm 1.

As discussed before, in order for the algorithm to converge, we must have $\lambda \succ \mathbf{0}$ in every iteration. We show that this is always the case. First, notice that we solve Problem (12) to update λ . Since $\left(\frac{1}{\zeta} \nabla_{\lambda} \Phi(\mathbf{u}_k) + 2\lambda_{\mathbf{u}_k} \right) \succ \mathbf{0}$, if $\lambda_{\mathbf{u}_k} \succ \mathbf{0}$, the optimum solution satisfies $\lambda_{k+1} \neq \mathbf{0}$. Therefore, if we initialize λ , such that $\lambda \succ \mathbf{0}$, it will hold that $\lambda_k \succ \mathbf{0}$ for all iterations. Secondly, it is not difficult to see from Lemma 1 that the optimal solution satisfies $\lambda^* \succ \mathbf{0}$ for $p < \infty$. Therefore, λ_k will safely converge to the optimum.

Note that, besides the algorithm we just introduced, when $p = 1$ we can optimize the model using block-coordinate descent as an alternative. In this case, the framework reduces to the traditional MT-MKL approach:

$$\min_{\mathbf{x} \in \Omega(\mathbf{x})} \mathbf{g}'(\mathbf{x}) \mathbf{1} \quad (17)$$

We can optimize with respect to $\{\mathbf{f}, \boldsymbol{\xi}, \mathbf{b}\}$ as a group, which involves T SVM problems, and then with respect to θ . The later parameter can be solved for via a closed-form expression.

Algorithm 1 Algorithm when $p \geq 1$

Choose M kernel functions. Calculate the kernel matrices \mathbf{K}_m^t for the T tasks and the M kernels. Initialize ζ , $\mathbf{u}_0 \in \Omega(\mathbf{u})$, ϵ , and $k = 0$.
while The duality gap is larger than ϵ **do**
 Given \mathbf{u}_k , solve Problem (10) and get \mathbf{v}_k ;
 Solve the minimization problem in (11);
 if The inequality (11) is not satisfied **then**
 $\zeta \leftarrow \frac{\zeta}{2}$;
 else
 Set \mathbf{u}_{k+1} as the optimal of (11);
 $k \leftarrow k + 1$;
 end if
end while

4 Learning in the Non-Convex Case

In this section we provide a simple algorithm to solve our framework in the case, when $p \in (0, 1)$, which renders Problem (1) to be non-convex. In what follows, we cast Problem (3) to an equivalent problem, which can be optimized via a simple group-coordinate descent algorithm. We first state the following lemma that is given in [1].

Lemma 2. For $p \in [\frac{1}{2}, 1)$, $\boldsymbol{\lambda} \in \mathbb{R}^T$, $\mathbf{g} \in \mathbb{R}^T$, $\mathbf{g} \succeq \mathbf{0}$ and $\mathbf{g} \neq \mathbf{0}$, we have

$$\min_{\boldsymbol{\lambda} \in \bar{B}_{\boldsymbol{\lambda}, q}} \mathbf{g}' \boldsymbol{\lambda}^{-1} = \nu(\mathbf{g})_p \quad (18)$$

with $q \triangleq \frac{p}{1-p}$, and the optimal $\boldsymbol{\lambda}^*$ can be calculated as follows:

$$\boldsymbol{\lambda}^* = \left(\frac{\mathbf{g}}{\nu(\mathbf{g})_p} \right)^{1-p} \quad (19)$$

Note that as $p \in [\frac{1}{2}, 1)$, the minimization problem in Problem (18) is convex with respect to $\boldsymbol{\lambda}$. In the following lemma it is shown that, for $p \in (0, \frac{1}{2})$, a similar convex equivalency can be constructed.

Lemma 3. For $p \in (0, \frac{1}{2})$, $\boldsymbol{\phi} \in \mathbb{R}^T$, $\mathbf{g} \in \mathbb{R}^T$, $\mathbf{g} \succeq \mathbf{0}$ and $\mathbf{g} \neq \mathbf{0}$, we have

$$\min_{\boldsymbol{\phi} \in \bar{B}_{\boldsymbol{\phi}, 1}} \mathbf{g}' \boldsymbol{\phi}^{-\frac{1}{q}} = \nu(\mathbf{g})_p \quad (20)$$

with $q \triangleq \frac{p}{1-p}$, and the optimal $\boldsymbol{\phi}^*$ can be calculated as follows:

$$\boldsymbol{\phi}^* = \left(\frac{\mathbf{g}}{\nu(\mathbf{g})_p} \right)^p \quad (21)$$

Note that Problem (20) is convex for $p \in (0, \frac{1}{2})$. The proof of Lemma 3 is given in Section .4 of the Appendix. The next lemma illustrates that, as $p \in (0, \frac{1}{2})$, even though Problem (18) is not convex, it is equivalent to the convex optimization problem (20).

Lemma 4. Under the conditions of Lemma 3, we have that

$$\min_{\boldsymbol{\lambda} \in \bar{B}_{\boldsymbol{\lambda}, q}} \mathbf{g}' \boldsymbol{\lambda}^{-1} \quad (22)$$

is equivalent to Problem (20) with optimum solution

$$\boldsymbol{\lambda}^* = \left(\frac{\mathbf{g}}{\nu(\mathbf{g})_p} \right)^{1-p} \quad (23)$$

The above lemma can be simply proved by letting $\phi = \lambda^p$. It states the fact that Equation (18) and Equation (19) are true not only in the case for $p \in [\frac{1}{2}, 1)$, when Problem (18) is convex, but also for $p \in (0, \frac{1}{2})$, when Problem (18) is not convex. This fact directly leads to the following theorem:

Theorem 3. Let $q \triangleq \frac{p}{1-p}$. For $p \in (0, 1)$, MT-MKL optimization Problem (1) is equivalent to the following problem

$$\min_{\mathbf{x} \in \Omega(\mathbf{x}), \boldsymbol{\lambda} \in \bar{B}_{\boldsymbol{\lambda}, q}} \mathbf{g}'(\mathbf{x}) \boldsymbol{\lambda}^{-1} \quad (24)$$

In order to optimize Problem (1), we can equivalently optimize Problem (24), where a simple group-coordinate descent method can be applied. In each iteration, given $\boldsymbol{\theta}$ and $\boldsymbol{\lambda}$, we first optimize with respect to $\{\mathbf{f}, \boldsymbol{\xi}, \mathbf{b}\}$. This problem involves T independent SVM problems, which can be solved using existing efficient SVM solvers, such as LIBSVM [5]. Then, the optimization problem with respect to $\boldsymbol{\theta}$ becomes:

$$\min_{\boldsymbol{\theta} \in \bar{B}_{\boldsymbol{\theta}, s}} \sum_{m=1}^M \frac{1}{\theta_m} \sum_{t=1}^T \frac{\|f_m^t\|_{H_m}^2}{\lambda_t} \quad (25)$$

Note that $\boldsymbol{\lambda} \succeq \mathbf{0}$ is always satisfied due to the constraint on $\boldsymbol{\lambda}$. Therefore, the aforementioned problem can be readily solved by applying Lemma 2, which supplies a closed-form solution for $\boldsymbol{\theta}$. Finally, Equation (19) provides a closed-form expression to update $\boldsymbol{\lambda}$.

The previous group-coordinate descent algorithm is justified as follows. Assume $s > 1$, as $p \in [\frac{1}{2}, 1)$, Problem (24) is convex with respect to each of the three blocks of variables, namely, $\{\mathbf{f}, \boldsymbol{\xi}, \mathbf{b}\}$, $\boldsymbol{\theta}$ and $\boldsymbol{\lambda}$. The convergence of the algorithm is therefore guaranteed in this case, based on [3, Prop. 2.7.1]. This is majorly because the SVM solver provides a unique solution for $\{\mathbf{f}, \boldsymbol{\xi}, \mathbf{b}\}$, and the solutions for $\boldsymbol{\theta}$ and $\boldsymbol{\lambda}$ are also uniquely obtained based on Lemma 2. For a detailed proof, we refer the readers to the proof of Theorem 4 in [18], which is similar in spirit. When $p \in (0, \frac{1}{2})$, optimizing with respect to $\{\mathbf{f}, \boldsymbol{\xi}, \mathbf{b}\}$ and $\boldsymbol{\theta}$ is the same as the case for $p \in [\frac{1}{2}, 1)$. The only difference is that, Problem (24) is not convex with respect to $\boldsymbol{\lambda}$. However, due to Lemma 4, minimizing it with respect to $\boldsymbol{\lambda}$ is equivalent to the convex problem (20). Hence, similar to the case when $p \in [\frac{1}{2}, 1)$, we iteratively solve three convex problems with respect to $\{\mathbf{f}, \boldsymbol{\xi}, \mathbf{b}\}$, $\boldsymbol{\theta}$ and $\boldsymbol{\lambda}$ respectively. The algorithm will be able to converge, again, because of [3, Prop. 2.7.1] for the same reason as previously stated. The algorithm is summarized in Algorithm 2.

Algorithm 2 Algorithm when $p \in (0, 1)$

Choose M kernel functions. Calculate the kernel matrices \mathbf{K}_m^t for the T tasks and the M kernels. Initialize $\boldsymbol{\theta}_0 \in \bar{B}_{\boldsymbol{\theta}, s}$, $\boldsymbol{\lambda}_0 \in \bar{B}_{\boldsymbol{\lambda}, q}$, ϵ , and $k = 0$.

while The change of objective function value is larger than ϵ **do**

 Given $\boldsymbol{\theta}_k$, $\boldsymbol{\lambda}_k$, solve T SVM problems, get \mathbf{f}_k ;

 Given \mathbf{f}_k , $\boldsymbol{\lambda}_k$, calculate $\boldsymbol{\theta}_{k+1}$ based on Lemma 2;

 Given \mathbf{f}_k , $\boldsymbol{\theta}_{k+1}$, calculate $\boldsymbol{\lambda}_{k+1}$ based on Equation (19);

$k \leftarrow k + 1$;

end while

It is worth emphasizing that our proposed algorithm for $p < 1$ has the same asymptotic computational complexity compared to the algorithm for $p = 1$, which was introduced at the end of Section 3. In each iteration, the algorithm for $p = 1$ involves T SVM optimizations and a closed-form computation for $\boldsymbol{\theta}$, while the algorithm for $p < 1$ adds only one more closed-form computation for $\boldsymbol{\lambda}$. Hence, the computational complexity in each iteration is dominated by the one of the SVM optimizer. On the other hand, the algorithm for $p \geq 1$ involves solving a quadratic programming problem (12) and a closed-form calculation for solving Problem (13). Therefore, in each iteration, the computational complexity is dominated by the one of the Problem (12) solver.

5 Experiments

In order to assess the merits of the new framework, we experimented with a few, selected data sets. In this section we present the obtained results and discuss the effects of varying p .

5.1 Experimental Settings

Our experiments involve 6 multi-class problems, namely, Wall-Following Robot Navigation (*Robot*), Image Segmentation (*Segment*), Landsat Satellite (*Satellite*), Statlog Shuttle (*Shuttle*), Vehicle Silhouettes (*Vehicle*) and Steel Plates Faults (*Steel*) data set. All data sets were retrieved from the UCI Machine Learning Repository [8]. For all data sets, an equal number of samples from each class was chosen for training. Note that the original *Shuttle* data set has seven classes, four of which are poorly represented; for this data set we only chose data from the other three classes. The attributes of all data sets were normalized so that they lie in $[0, 1]$. Finally, we modeled each multi-class problem as an MTL problem, whose recognition tasks considered every possible combination of a class versus another class (*i.e.* $\binom{c}{2}$ one-vs-one tasks for c classes).

Besides the UCI benchmark problems, we also performed experiments on two widely used multi-task data set, namely the *Letter* and *Landmine* data sets, which are introduced in detail later. For all experiments, 11 kernels were pre-specified: the Linear kernel, 2^{nd} -order Polynomial kernel and Gaussian kernels with spread parameter values $\{2^{-7}, 2^{-5}, 2^{-3}, 2^{-1}, 2^0, 2^1, 2^3, 2^5, 2^7\}$. The value of SVM's parameter C was selected via cross-validation. Also, the value of parameter s , appearing in the norm constraint of θ , was held fixed to 1.1. Finally, in order to discuss the effect of p , we allowed it to take values in $\{0.01, 0.02, 0.05, 0.1, 0.2, 0.5, 1, 2, 5, 10, 20, 50, \infty\}$.

5.2 Effect of p on Objective Function Values

In Figure 1, we first show how the optimum objective function values of the T tasks change as p varies. In our subsequent discussion we'll focus on the results for the *Robot* data set, due to its small number of tasks and, hence, is easy to analyze; similar observations can be made for the remaining data sets.




Figure 1: Experimental results for the *Robot* data set. Objective function value of each task as a function of p .

Observations: The optimum objective value of task 2, which is the one exhibiting the highest objective value among the 6 tasks, decreases, when p increases, and achieves its lowest value for $p \geq 5$. Similar behavior is observed for task 4, which is the task that has the second highest objective value. On the contrary, the other tasks have growing objective values as p increases. Also, according to Proposition 1, every p value

yields a PF point of the relevant T -objective MOO problem. This can be observed in the figure, since there is no p such that the corresponding T objective values lead to a dominating solution.

Discussion: The behavior displayed in Figure 1 can be explained as follows. When $p \geq 1$, according to Lemma 1, we know that $\lambda^* \propto g^{p-1}$. Therefore, the λ_t 's corresponding to tasks with high objective values are larger than the remaining λ_t 's. This means that these tasks are more heavily penalized as p increases. In the extreme case, where $p \rightarrow \infty$, task t_0 , that has the highest objective value, will have $\lambda_{t_0} = 1$ and the other λ_t 's will all be zero. This amounts to only task t_0 being penalized (thus optimized), while the other tasks' performances are ignored. Similarly, when $p < 1$, Lemma 2 and Lemma 4 imply that $\lambda^* \propto g^{1-p}$. Therefore, in this case, the smaller p is, the heavier the tasks with low objective values are penalized. This explains the trend of each curve in Figure 1. For the other data sets, similar observations can be stated for the same reasons.

5.3 Effect of p on Classification Performances

We now discuss how p 's value affects the classification performance of all tasks. Again, we analyze the performance on the *Robot* data set in detail. Figure 2 depicts the correct classification rate for each of the 6 binary classification tasks, while Figure 3 illustrates the overall classification rate as a multi-class problem. All experiments were performed using 5% of the data available for training and the averages over 20 runs are reported. Note that Figure 2 showcases how the *difference in classification rate* (DCR) changes with p . Here, DCR is defined as the difference in correct classification rates for each p with respect to the rate when $p = 1$. The latter rate is displayed in parenthesis inside the legend for each task. As mentioned earlier, the latter case corresponds to the traditional MT-MKL approach of optimizing the average of objective function values (all objectives are equally weighted).




Figure 2: Experimental results for the *Robot* data set. Difference in Classification Rate (DCR) for each task with changing p .

Observations: Upon inspection of Figure 2, one can immediately observe that the correct classification rate increases as p decreases for each curve. Moreover, the best result is always achieved when $p < 1$ for all tasks. It can be seen that tasks 2 and 4, achieve the most significant improvement in classification rate, when $p < 1$. On the other hand, the other four tasks also enjoy improved performance for $p < 1$, when compared to $p = 1$. We used a t-test with significance level $\alpha = 0.05$ to compare the improvement between $p = 1$ and $p = 0.01$. The results of these tests confirmed the statistically significant improvement for all tasks except task 3. Figure 3 demonstrates the same behavior as the one shown in Figure 2. Note that the red dashed line shows the accuracy obtained when $p = 1$. We immediately observe that the best performance is achieved as p decreases towards 0.01. Again, t-tests show statistically significant improvements.




Figure 3: Experimental results for the *Robot* data set. The black solid curve shows the overall multi-class classification accuracy as p varies, while the red dashed line depicts the accuracy obtained, when $p = 1$.

Discussion: The above observations can be explained as follows. The aim of the proposed MT-MKL framework is to find an RKHS by linearly combining pre-selected kernels via the coefficients θ , such that the T tasks can achieve good performance. By applying Lemma 1 to Problem (9) and Problem (24) for θ , one can easily see that the optimum solution is such that $\theta_m \propto (\sum_{t=1}^T \eta^t G_m^t)^{\frac{1}{s-1}}$, where we define $G_m^t \triangleq \alpha^{t'} \mathbf{Y}^{t'} \mathbf{K}_m^t \mathbf{Y}^t \alpha^t$ and $\eta^t \triangleq \lambda^t$, when $p \geq 1$, and $\eta^t \triangleq \frac{1}{\lambda^t}$, when $p < 1$. Also, as demonstrated earlier, we know that $\eta^t \propto (g^t)^{p-1}$ for $p > 0$.

Based on these facts, let us first discuss the behavior as $p \rightarrow \infty$. In this scenario, as previously mentioned, the model is only optimizing task t_0 , *i.e.* the one with the highest objective value. It will hold that $\eta_{t_0} = 1$, while the other η_t 's will be 0. In this case, the θ_m 's are only determined by $G_m^{t_0}$. In other words, the RKHS is learned only by the training samples from the t_0 -th task, which is a (potentially, very) small proportion of the union of task-specific training sets. Hence, it is unsurprising that the correct classification rate is so low in this case.

When p is finite but very high, the η_t 's that correspond to tasks with high objective values will be much higher than the other η_t 's. As p decreases, the former η_t 's will decrease in value. On the other hand, the other tasks will have increased η_t 's. This means that the θ_m 's are estimated not only based on the training data for the tasks with high objective values, but also based on the training samples from other tasks. Therefore the classification accuracy will be improving for each task. This behavior continues as p decreases towards 1, until all η_t 's are equal. However, even though θ_m is now determined by the average of G_m^t 's, it does not necessarily mean that the training data of each task have the same influence on estimating the θ_m 's. The reason is that the tasks with high objective values usually have much larger G_m^t compared to the other tasks due to smaller SVM margins (note that $G_m^t = \|f_m^t\|_{H_m}^2$). Thus, even though θ_m is determined by the average of G_m^t 's, the tasks with high objective values will influence this average value the most. One can interpret this as θ_m being estimated primarily based on the training data of these high objective valued tasks and to a lesser degree based on the training data of the other tasks.

Considering these last remarks, it is now straightforward to explain, why the classification performances obtained via the use of the proposed MT-MKL framework is better for $p < 1$. In this case, the η^t 's have larger values for tasks with lower objective values. Therefore, the G_m^t 's associated to the these tasks perform a role of ever increasing importance as p decreases, until some threshold value of p , after which the $\eta^t G_m^t$'s have similar values for all t . This can be interpreted as the θ_m 's being estimated by considering the whole training set from all T tasks, which eventually leads to improved performance for all tasks involved.

In Table 1, we show the overall classification performance results for all data sets with different size of

Table 1: Comparison of multi-class classification accuracy between $p < 1$, $p = 1$ and $p \rightarrow \infty$

10%	Robot	Sat	Vec	Steel	Seg	Shuttle
$p < 1$	90.11	85.48	61.47	73.35	90.22	98.77
$p = 1$	<u>88.63</u>	<u>83.79</u>	60.61	73.33	<u>88.56</u>	<u>97.78</u>
$p \rightarrow \infty$	87.01	82.12	46.75	71.49	84.98	72.32
# tasks	6	15	6	3	21	3
# tasks \uparrow	6	12	5	2	21	1
20%	Robot	Sat	Vec	Steel	Seg	Shuttle
$p < 1$	92.70	87.78	66.97	76.33	92.86	99.20
$p = 1$	92.36	<u>85.43</u>	65.83	75.82	<u>90.72</u>	<u>98.74</u>
$p \rightarrow \infty$	91.51	83.55	50.80	72.74	87.71	76.03
# tasks	6	15	6	3	21	3
# tasks \uparrow	6	12	5	3	19	1
50%	Robot	Sat	Vec	Steel	Seg	Shuttle
$p < 1$	96.12	90.56	73.07	79.30	94.76	99.28
$p = 1$	96.03	<u>87.53</u>	<u>70.05</u>	<u>76.78</u>	<u>92.66</u>	<u>98.71</u>
$p \rightarrow \infty$	95.64	85.60	52.70	72.93	90.36	80.71
# tasks	6	15	6	3	21	3
# tasks \uparrow	6	14	6	3	21	2

training set. The percentage of the available samples that are used for training are shown in the first column for each set of experiments (10%, 20%, and 50% of the training set). The trends in overall accuracy are similar to the ones of task-wise performances, when p varies. As was the case with the *Robot* data set, the classification accuracy always decreases as p increases. Here, we only show three results for each data set, namely, the highest accuracy obtained when $p < 1$, the accuracy when $p = 1$ and the accuracy when $p \rightarrow \infty$. Since it is important to compare the performance between the cases of $p < 1$ and $p = 1$ (the traditional MT-MKL approach), a t-test was employed to test the statistical significance of the accuracy improvements between the cases of $p < 1$ and $p = 1$. In the table, underlined numbers indicate the results that are statistically significantly worse than the corresponding ones when $p < 1$. It can be seen that the results are consistent with our motivation and that they experimentally validate our analysis: using $p < 1$ always achieves significantly higher accuracy compared to the traditional MTL approach. Rows labeled as “# tasks” indicate the total number of tasks for each data set and rows labeled as “# tasks \uparrow ” show the number of tasks that experience an increase in classification rate as p decreases. It can be seen that for all data sets, most tasks improve in performance with decreasing p , which means that the best performance for these tasks is obtained when p is close to 0.

5.4 Comparison to Other Approaches

We also compared our framework to a baseline approach and three popular MTL models, namely *Sparse MTL* [24], *Tang’s method* [30], and *Minimax MTL* [20]. The baseline method trains each SVM task individually using MKL, *i.e.*, for each task t , the task-specific kernel coefficients $\theta^t \triangleq [\theta_1^t, \dots, \theta_M^t]'$ are learned using the constraint $\|\theta^t\|_q \leq 1$, $q \geq 1, \forall t$. Next, *Sparse MTL* [24] is a popular MTL model, which utilizes the $L_p - L_q$ mixed-norm $\sum_{m=1}^M (\sum_{t=1}^T \|f_m^t\|^q)^{p/q}$ as a regularizer. By letting $0 < p \leq 1$, this choice aims to encourage sparser kernel representation across tasks than the case when $p = 1$. Following the setting in [24], in our experiments we use the ranges $0 < p \leq 1$ and $1 \leq q \leq 2$ to identify the parameters p and q via cross-validation.

Next, *Tang’s method* [30] employs a different information sharing strategy. Unlike *Sparse MTL*, which relates tasks via the use of the mixed-norm, *Tang’s method* utilizes a partially shared kernel function for the same purpose. Specifically, the kernel function for each task is defined as $k^t \triangleq \sum_{m=1}^M (\xi_m + \gamma_m^t) k_m$ with

constraints $\sum_{m=1}^M (\xi_m + \gamma_m^t) = 1, \forall t; \xi_m \geq 0, \gamma_m^t \geq 0, \forall m, t; \sum_{m=1}^M \sum_{t=1}^T \gamma_m^t \leq \beta$. Obviously, by letting $\gamma_m^t = 0, \forall m, t$, the kernel function becomes the same for all tasks, which is the setting that we used in our method. Similar to several other MTL formulations, *Tang's method* minimizes the average loss of the T tasks. In our experiments, the parameter β is selected via cross-validation.

The recently proposed *Minimax MTL* [20] aims to guarantee a minimum level of accuracy for each task. Its motivation differs to that of our method, which considers improving the MTL performance by traversing the Pareto path of the associated MOO problem. Therefore, although similar, the two models are differently formulated: *Minimax MTL* minimizes the maximum (or more generally, the L_p -norm) of the T loss functions, while the L_p -norm of our model is applied to the T objectives (both regularizer and loss function). In our experiments, we followed the experimental settings of [20]: a linear model, *i.e.*, $f^t(\mathbf{x}) \triangleq \mathbf{w}^{t'} \mathbf{x}$, in conjunction with the hinge loss were utilized. Also, we used the model parameter values $\alpha = 0.1T$ and $0.2T$ (please refer to [20] for model details), and we report here the best results obtained. Two different regularizers were considered. The first regularizer was $\|\mathbf{W}\|_{tr}$, where $\mathbf{W} \triangleq [\mathbf{w}^1, \dots, \mathbf{w}^T]$ and $\|\mathbf{W}\|_{tr}$ denotes the trace norm, *i.e.*, the sum of the singular values of \mathbf{W} , as proposed by Argyriou, Evgeniou and Pontil (AEP) [2]. The second regularizer was $\|\mathbf{v}^0\|^2 + \frac{\lambda}{T} \sum_{t=1}^T \|\mathbf{v}^t\|^2$, where it is assumed that \mathbf{w}^t is given as $\mathbf{w}^t = \mathbf{v}^0 + \mathbf{v}^t, \forall t$ and the problem is optimized over \mathbf{v}^0 and the \mathbf{v}^t 's. This regularizer was proposed by Evgeniou and Pontil (EP) [7].

Table 2: Comparison of multi-class classification accuracy between our method and four other models

10%	Robot	Sat	Vec	Steel	Seg	Shuttle
Pareto	90.11	85.48	61.47	73.35	90.22	98.77
Sparse	<u>87.59</u>	85.43	60.58	75.21	90.01	98.66
Tang	<u>88.61</u>	85.27	60.50	73.03	<u>88.65</u>	98.06
MM-AEP	92.78	<u>83.06</u>	71.89	<u>65.65</u>	89.27	<u>97.22</u>
MM-EP	93.96	<u>80.94</u>	68.31	<u>69.45</u>	<u>85.57</u>	<u>97.43</u>
Baseline	<u>87.93</u>	<u>82.46</u>	60.05	72.98	<u>83.48</u>	<u>74.81</u>
20%	Robot	Sat	Vec	Steel	Seg	Shuttle
Pareto	92.70	87.78	66.97	76.33	92.86	99.20
Sparse	<u>91.81</u>	87.99	67.58	76.44	92.77	99.08
Tang	92.13	87.40	66.67	75.07	<u>91.17</u>	<u>98.83</u>
MM-AEP	95.10	<u>84.63</u>	75.07	<u>66.83</u>	<u>91.30</u>	<u>97.56</u>
MM-EP	96.59	<u>82.08</u>	73.46	<u>72.45</u>	<u>89.89</u>	<u>97.72</u>
Baseline	<u>91.45</u>	<u>84.15</u>	65.61	75.17	<u>85.30</u>	<u>78.94</u>
50%	Robot	Sat	Vec	Steel	Seg	Shuttle
Pareto	96.12	90.56	73.07	79.30	94.76	99.28
Sparse	95.77	92.18	74.85	78.84	94.49	99.66
Tang	95.68	90.52	74.04	78.64	94.59	99.43
MM-AEP	97.71	<u>86.85</u>	78.80	<u>67.54</u>	<u>92.55</u>	<u>97.98</u>
MM-EP	98.73	<u>85.49</u>	75.57	<u>75.26</u>	<u>93.15</u>	<u>98.37</u>
Baseline	<u>95.09</u>	<u>86.32</u>	<u>69.89</u>	<u>76.66</u>	<u>86.32</u>	<u>82.59</u>

The experimental results are reported in Table 2, where we compare our method, labeled *Pareto* on the first row of each section of the table, and the four aforementioned methods. Note that for our method, we reported the best result achieved when $p < 1$. For the *Minimax MTL* method, we reported results, when both the AEP and EP regularizers are applied (labeled *MM-AEP* and *MM-EP* respectively). Underlined values represent results that are worse than the ones of our method in a statistically significant sense.

First, by comparing the MTL methods to the baseline, we observe that, utilizing MTL is indeed beneficial in improving the classification performance, with only very few exceptions. In addition, our method always outperforms the baseline model. These facts validate the effectiveness of the MTL methodology. Secondly, we observe that, in most experiments, our method outperforms the other models and many of the observed

improvements are statistically significant. In other words, in practice, there are indeed multi-task classification problems that can benefit from our model, for which the L_p -norm, instead of the average, of the T objectives is minimized.

In the sequel, we report results on two multi-task data sets mentioned earlier in the text. The first one, namely the *Letter* data set¹, is a collection of handwritten words compiled by Rob Kassel of the MIT Spoken Language Systems Group. The associated MTL problem involves 8 tasks, each of which is a binary classification problem for handwritten letters. The 8 tasks are: ‘C’ vs. ‘E’, ‘G’ vs. ‘Y’, ‘M’ vs. ‘N’, ‘A’ vs. ‘G’, ‘I’ vs. ‘J’, ‘A’ vs. ‘O’, ‘F’ vs. ‘T’ and ‘H’ vs. ‘N’. Each letter is represented by a 8×16 pixel image, which forms a 128-dimensional feature vector. The goal for this problem is to correctly recognize the letters in each task. For the experiments, due to the large size of the data set (45679 data in total), we randomly chose 200 samples for each task.

On the other hand, the *Landmine* data set² consists of 29 binary classification tasks. Each datum is a 9-dimensional feature vector extracted from radar images that capture a single region of landmine fields. The 9 features include four moment-based features, three correlation-based features, one energy ratio feature and one spatial variance feature [36]. Tasks 1 – 15 correspond to regions that are relatively highly foliated, while the other 14 tasks correspond to regions that are bare earth or desert. The tasks entail different amounts of data, varying from 30 to 96 samples. The goal is to detect landmines in specific regions.

Table 3: Comparison of multi-task classification accuracy between $p < 1$, $p = 1$ and $p \rightarrow \infty$

Letter	10%	20%	50%
$p < 1$	83.95	87.51	90.61
$p = 1$	<u>81.45</u>	<u>86.42</u>	90.01
$p \rightarrow \infty$	77.67	85.60	89.80
# tasks	8	8	8
# tasks \uparrow	8	8	7
Landmine	20%	30%	50%
$p < 1$	69.59	74.15	77.42
$p = 1$	<u>67.24</u>	<u>71.62</u>	76.96
$p \rightarrow \infty$	65.01	69.00	74.06
# tasks	29	29	29
# tasks \uparrow	29	28	27

The experimental settings for these two problems were the same as the ones used in the previous experiments, except that we did not use 10% of the *Landmine* data set for training, due to the small size of the data set; instead, we used 20%, 30% and 50% of the available training data. The averages of classification accuracy of each task are displayed in Table 3.

We can observe from Table 3 and Table 4 that the experimental results on the two benchmark multi-task problems are, again, consistent with our previous analysis. It is worth noting that the improved performance for $p < 1$ is achieved without incurring additional computational cost, since, as stated in Section 4, the algorithm for $p < 1$ has the same asymptotic computational complexity per iteration as the one for $p = 1$.

6 Conclusions

To recapitulate, this paper introduced a new Support Vector Machine (SVM)-based Multi-Task Multiple Kernel Learning (MT-MKL) framework that co-trains binary classification tasks. The task objectives are forced to share a common kernel, which is inferred from the data via Multiple Kernel Learning (MKL). Unlike traditional Multi-Task Learning (MTL) approaches, which only consider optimizing the average of objective functions, it can be shown that our framework optimizes a specific variety of (implicitly-defined)

¹Available at: <http://multitask.cs.berkeley.edu/>

²Available at: <http://people.ee.duke.edu/~lcarin/LandmineData.zip>

Table 4: Comparison of multi-task classification accuracy between our method and four other models

Letter	10%	20%	50%
Pareto	83.95	87.51	90.61
Sparse	83.00	87.09	90.65
Tang	<u>80.86</u>	<u>82.95</u>	<u>84.87</u>
MM-AEP	83.92	86.77	<u>88.94</u>
MM-EP	<u>81.73</u>	<u>84.14</u>	<u>86.43</u>
Baseline	<u>81.33</u>	86.39	<u>89.80</u>
Landmine	20%	30%	50%
Pareto	69.59	74.15	77.42
Sparse	<u>58.89</u>	<u>65.83</u>	<u>75.82</u>
Tang	<u>66.60</u>	<u>70.89</u>	76.08
MM-AEP	<u>59.27</u>	<u>61.06</u>	<u>67.34</u>
MM-EP	69.69	73.62	76.40
Baseline	<u>66.64</u>	<u>71.14</u>	76.29

conic combinations of these functions. These combinations are parameterized by a parameter $p > 0$ and specialize to the average of task objectives, when $p = 1$.

Our motivation to construct such a framework stemmed from the observation that the solution obtained by optimizing the average of objective functions coincides with a particular point (solution) on the Pareto Front (PF) (the set of all non-dominated solutions) of a Multi-Objective Optimization (MOO) problem, which considers the simultaneous vector-optimization of all objective functions involved in the MTL setting. In the case of our framework, we showed that the conic combinations it considers trace out a path on the PF by varying p . This path embodies the trade-off in perceived importance among task objectives rather than among task-related classification performances. Since the traditional MTL approach is largely a heuristic, we argued that, by considering other conic combinations (thus, other points on the aforementioned path), we may actually be able to uniformly improve the classification performances of all tasks even more.

In order to generate solutions for the framework’s formulation, through a series of equivalence results we had to distinguish between two cases: $p \geq 1$ (convex case) and $p \in (0, 1)$ (non-convex case). For each of these two cases we derived algorithms, which are straightforward to implement, as they can leverage from existing solvers and closed-form solutions, which we derived.

Towards assessing the framework’s merits, we performed a series of experiments on 6 benchmark multi-class recognition problems, which were modeled as multi-task problems suitable to our framework. Additionally, 2 multi-task data sets were also considered. The obtained experimental results demonstrated clear advantages of considering only the range $0 < p < 1$ for all data sets. When compared to the traditional MTL approach ($p = 1$), these advantages were reflected in the simultaneous classification improvements for each task individually, as well as in the improvement of overall accuracy for the multi-class recognition problems we considered. We justified this phenomenon as follows: when p decreases towards 0, the objectives with lower value become more important, which implicitly increases the effective number of training samples for all tasks and allows each task-specific model to improve its classification accuracy.

Future works involve formulating Pareto-Path MT-MKL models based on other kernel machines, such as Kernel Ridge Regression [26], One-class SVM [27], Support Vector Domain Description [31], etc. It will be of interest to analyze the behavior of these models, when the Pareto-Path MT-MKL framework is adapted. Also, it would be of interest to derive generalization bounds for this new MT-MKL framework, so that its performance merits are better explained.

Acknowledgments

C. Li acknowledges partial support from National Science Foundation (NSF) grant No. 0806931 and No. 0963146. Moreover, M. Georgiopoulos acknowledges partial support from NSF grants No. 0963146, No. 1200566, and No. 1161228. Finally, G. C. Anagnostopoulos acknowledges partial support from NSF grants No. 0647018 and No. 1263011. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

References

- [1] Jonathan Aflalo, Aharon Ben-Tal, Chiranjib Bhattacharyya, Jagarlapudi Saketha Nath, and Sankaran Raman. Variable sparsity kernel learning. *Journal of Machine Learning Research*, 12:565–592, 2011.
- [2] Andreas Argyriou, Theodoros Evgeniou, and Massimiliano Pontil. Convex multi-task feature learning. *Machine Learning*, 73(3):243–272, 2008.
- [3] Dimitri P. Bertsekas. *Nonlinear Programming: 2nd Edition*. Athena Scientific, 1999.
- [4] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [5] Chih-Chung Chang and Chih-Jen Lin. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27, 2011. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- [6] Corinna Cortes, Mehryar Mohri, and Afshin Rostamizadeh. Learning non-linear combinations of kernels. In *Advances in Neural Information Processing Systems*, pages 396–404, 2009.
- [7] Theodoros Evgeniou and Massimiliano Pontil. Regularized multi-task learning. In *Proceedings of the tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 109–117, 2004.
- [8] A. Frank and A. Asuncion. UCI machine learning repository, 2010. Available from: <http://archive.ics.uci.edu/ml>.
- [9] Kun Gai, Guangyun Chen, and Chang-shui Zhang. Learning kernels with radiuses of minimum enclosing balls. In *Advances in Neural Information Processing Systems*, pages 649–657, 2010.
- [10] Mehmet Gönen and Ethem Alpaydin. Localized multiple kernel learning. In *Proceedings of the 25th International Conference on Machine Learning*, pages 352–359, 2008.
- [11] Mehmet Gonen and Ethem Alpaydin. Multiple kernel learning algorithms. *Journal of Machine Learning Research*, 12:2211–2268, 2011.
- [12] M. Grant and S. Boyd. Graph implementations for nonsmooth convex programs. In V. Blondel, S. Boyd, and H. Kimura, editors, *Recent Advances in Learning and Control*, Lecture Notes in Control and Information Sciences, pages 95–110. Springer-Verlag Limited, 2008. http://stanford.edu/~boyd/graph_dcp.html.
- [13] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 1.21, April 2011.
- [14] Pratik Jawanpuria and J. Saketha Nath. Multi-task multiple kernel learning. In *2011 SIAM International Conference on Data Mining*, 2011.
- [15] Tony Jebara. Multi-task feature and kernel selection for SVMs. In *the 21st International Conference on Machine Learning*, 2004.
- [16] Marius Kloft and Gilles Blanchard. On the convergence rate of l_p -norm multiple kernel learning. *Journal of Machine Learning Research*, 13:2465–2501, 2012.

- [17] Marius Kloft, Ulf Brefeld, Pavel Laskov, and Soren Sonnenburg. Non-sparse multiple kernel learning. In *NIPS Workshop on Kernel Learning: Automatic Selection of Optimal Kernels*, 2008.
- [18] Marius Kloft, Ulf Brefeld, Soren Sonnenburg, and Alexander Zien. l_p -norm multiple kernel learning. *Journal of Machine Learning Research*, 12:953–997, 2011.
- [19] Gert R.G. Lanckriet, Nello Cristianini, Peter Bartlett, Laurent El Ghaoui, and Michael I. Jordan. Learning the kernel matrix with semidefinite programming. *Journal of Machine Learning Research*, 5:27–72, 2004.
- [20] Nishant Mehta, Dongryeol Lee, and Alexander Gray. Minimax multi-task learning and a generalized loss-compositional paradigm for MTL. In *Advances in Neural Information Processing Systems*, pages 2159–2167, 2012.
- [21] Arkadi Nemirovski. Prox-method with rate of convergence $o(1/t)$ for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal of Optimization*, 15:229–251, 2005.
- [22] Cheng Soon Ong, Alexander J. Smola, and Robert C. Williamson. Learning the kernel with hyperkernels. *Journal of Machine Learning Research*, 6:1043–1071, 2005.
- [23] Alain Rakotomamonjy, Francis R. Bach, Stephane Canu, and Yves Grandvalet. SimpleMKL. *Journal of Machine Learning Research*, 9:2491–2521, 2008.
- [24] Alain Rakotomamonjy, Remi Flamary, Gilles Gasso, and Stephane Canu. l_p-l_q penalty for sparse linear and sparse multiple kernel multitask learning. *IEEE Transactions on Neural Networks*, 22:1307–1320, 2011.
- [25] Wojciech Samek, Alexander Binder, and Motoaki Kawanabe. Multi-task learning via non-sparse multiple kernel learning. In Pedro Real, Daniel Diaz-Pernil, Helena Molina-Abril, Ainhoa Berciano, and Walter Kropatsch, editors, *Computer Analysis of Images and Patterns*, volume 6854 of *Lecture Notes in Computer Science*, pages 335–342. Springer Berlin / Heidelberg, 2011.
- [26] Craig Saunders, Alexander Gammerman, and Volodya Vovk. Ridge regression learning algorithm in dual variables. In *Proceedings of the 15th International Conference on Machine Learning*, pages 515–521, 1998.
- [27] Bernhard Schölkopf, John C. Platt, John Shawe-Taylor, and Alex J. Smola. Estimating the support of a high-dimensional distribution. *Neural Computation*, 13:1443–1471, 2001.
- [28] Maurice Sion. On general minimax theorems. *Pacific Journal of Mathematics*, 8:171–176, 1958.
- [29] Soren Sonnenburg, Gunnar Ratsch, Christin Schafer, and Bernhard Scholkopf. Large scale multiple kernel learning. *Journal of Machine Learning Research*, 7:1531–1565, 2006.
- [30] Lei Tang, Jianhui Chen, and Jieping Ye. On multiple kernel learning with multiple labels. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, pages 1255–1260, 2009.
- [31] David M.J. Tax and Robert P.W. Duin. Support vector domain description. *Pattern Recognition Letters*, 20:1191–1199, 1999.
- [32] Paul Tseng. On accelerated proximal gradient methods for convex-concave optimization. *Submitted to SIAM Journal on Optimization*, 2008.
- [33] Manik Varma and Bodla Rakesh Babu. More generality in efficient multiple kernel learning. In *Proceedings of the 26th International Conference on Machine Learning*, pages 1065–1072, 2009.
- [34] Christian Widmer, Nora C. Toussaint, Yasemin Altun, and Gunnar Ratsch. Multitask multiple kernel learning (mt-mkl). In *NIPS 2010 Workshop: New Directions in Multiple Kernel Learning*, 2010.

- [35] Zenglin Xu, Rong Jin, Haiqin Yang, Irwin King, and Michael R Lyu. Simple and efficient multiple kernel learning by group lasso. In *Proceedings of the 27th International Conference on Machine Learning*, pages 1175–1182, 2010.
- [36] Ya Xue, Xuejun Liao, Laurence Carin, and Balaji Krishnapuram. Multi-task learning for classification with dirichlet process priors. *Journal of Machine Learning Research*, 8:35–63, 2007.

.1 Proof of Proposition 1

Assume that \mathbf{x}^* is not a Pareto Front solution. Then there is a feasible point \mathbf{x}^{**} satisfying $g_t(\mathbf{x}^{**}) \leq g_t(\mathbf{x}^*)$, $t \in Z_T$, while at least for one t it holds that $g_t(\mathbf{x}^{**}) < g_t(\mathbf{x}^*)$. This means that $\nu(\mathbf{g}(\mathbf{x}^{**}))_p < \nu(\mathbf{g}(\mathbf{x}^*))_p$, which contradicts our assumption that \mathbf{x}^* is the optimal of Problem (4).

.2 Proof of Theorem 1

Due to Lemma 1, Problem (1) (or 3) is equivalent to

$$\min_{\mathbf{x} \in \Omega(\mathbf{x})} \max_{\boldsymbol{\lambda} \in \bar{B}_{\boldsymbol{\lambda},q}} \boldsymbol{\lambda}' \mathbf{g}(\mathbf{x}) \quad (26)$$

when $p \geq 1$, where $q \triangleq \frac{p}{p-1}$. Here we remind the readers that \mathbf{x} encompasses all the Support Vector Machine (SVM)-related variables $\mathbf{f}, \boldsymbol{\theta}, \boldsymbol{\xi}, \mathbf{b}$. Since the objective function above is jointly convex with respect to $\mathbf{f}, \boldsymbol{\theta}, \boldsymbol{\xi}, \mathbf{b}$ and concave with respect to $\boldsymbol{\lambda}$, while the feasible set of the above problem is convex and compact, the order of minimization and maximization can be interchanged based on the Sion-Kakutani *min-max* theorem [28], which yields:

$$\max_{\boldsymbol{\lambda} \in \bar{B}_{\boldsymbol{\lambda},q}} \min_{\mathbf{x} \in \Omega(\mathbf{x})} \boldsymbol{\lambda}' \mathbf{g}(\mathbf{x}) \quad (27)$$

The inner minimization problem can be partially converted to the dual form of the SVM problem:

$$\begin{aligned} \max_{\boldsymbol{\lambda}} \min_{\boldsymbol{\theta}} \max_{\boldsymbol{\alpha}} \quad & \tilde{\Phi}(\boldsymbol{\theta}, \boldsymbol{\alpha}, \boldsymbol{\lambda}) \\ \text{s.t.} \quad & \boldsymbol{\alpha}^{t'} \mathbf{y}^t = 0, \quad \mathbf{0} \preceq \boldsymbol{\alpha}^t \preceq C\mathbf{1}, \quad t \in N_T; \\ & \boldsymbol{\theta} \in \bar{B}_{\boldsymbol{\theta},s}, s \geq 1; \quad \boldsymbol{\lambda} \in \bar{B}_{\boldsymbol{\lambda},q} \end{aligned} \quad (28)$$

where $\tilde{\Phi}(\boldsymbol{\theta}, \boldsymbol{\alpha}, \boldsymbol{\lambda}) = \sum_{t=1}^T \lambda^t (\boldsymbol{\alpha}^{t'} \mathbf{1} - \frac{1}{2} \boldsymbol{\alpha}^{t'} \mathbf{Y}^{t'} (\sum_{m=1}^M \theta_m \mathbf{K}_m^t) \mathbf{Y}^t \boldsymbol{\alpha}^t)$, $\mathbf{y}^t \in \{-1, 1\}^{N_t}$ is the vector containing all labels of the training data for the t -th task, \mathbf{Y}^t is a diagonal matrix with the elements of \mathbf{y}^t on its diagonal, and \mathbf{K}_m^t is the kernel matrix with elements $k_m(\mathbf{x}_i^t, \mathbf{x}_j^t)$. Since the function $\max_{\boldsymbol{\alpha}} \tilde{\Phi}(\boldsymbol{\theta}, \boldsymbol{\alpha}, \boldsymbol{\lambda})$ and the feasible set of the above problem satisfy the Sion-Kakutani *min-max* theorem [28], the order of the outer maximization and the minimization can be changed once again, which results in the following optimization problem:

$$\begin{aligned} \min_{\boldsymbol{\theta}} \max_{\boldsymbol{\lambda}, \boldsymbol{\alpha}} \quad & \tilde{\Phi}(\boldsymbol{\theta}, \boldsymbol{\alpha}, \boldsymbol{\lambda}) \\ \text{s.t.} \quad & \boldsymbol{\alpha}^{t'} \mathbf{y}^t = 0, \quad \mathbf{0} \preceq \boldsymbol{\alpha}^t \preceq C\mathbf{1}, \quad t \in N_T; \\ & \boldsymbol{\theta} \in \bar{B}_{\boldsymbol{\theta},s}, s \geq 1; \quad \boldsymbol{\lambda} \in \bar{B}_{\boldsymbol{\lambda},q} \end{aligned} \quad (29)$$

Next, if we consider the change of variables $\beta_i^t \triangleq \alpha_i^t \lambda^t$, $t \in N_T$, then Problem (29) becomes Problem (9). Note that, in fact, the constraint $\boldsymbol{\beta}^{t'} \mathbf{y}^t = 0$ should have been $\frac{\boldsymbol{\beta}^{t'} \mathbf{y}^t}{\lambda^t} = 0$ if $\lambda^t \neq 0$ and $\boldsymbol{\beta}^t = \mathbf{0}$ if $\lambda^t = 0$. However, these two constraints are equivalent, even in the case, when $\lambda^t = 0$, as we must have $\boldsymbol{\beta}^t = \mathbf{0}$, because of the constraint $\boldsymbol{\beta}^t \preceq \lambda^t \mathbf{1}$.

.3 Proof of Theorem 2

The proof is separated into two parts. Initially, we prove the first part of the theorem, where $\Omega(\theta) \triangleq \{\theta \mid \theta \in \bar{B}_{\theta,s}\}$.

Notice that the minimization problem is convex, since Ω is a convex set and the objective function has a positive semi-definite Hessian matrix. Hence, it features a unique minimizer θ^* , which will be the unique stationary point (saddle point) of the problem's Lagrangian \mathcal{L} . Also, since $\{\theta \mid \nu(\theta)_s \leq 1\} = \{\theta \mid \nu(\theta)_s^s \leq 1\}$, we redefine $\Omega \triangleq \{\theta \mid \theta \succeq \mathbf{0}, \nu(\theta)_s^s \leq 1\}$. Notice that $(\theta)_s^s = \|\theta\|_s^s = \mathbf{1}'\theta^s$. Additionally, Problem (14) is equivalent to

$$\min_{\theta \in \Omega} [\mathbf{1}'\theta^s - s\psi'\theta] \quad (30)$$

after multiplying the original objective function by $s > 1$. If $\lambda \succeq \mathbf{0}$ and $\mu \geq 0$ are the dual variables associated to the constraints $\theta \succeq \mathbf{0}$ and $\|\theta\|_s^s \leq 1$ respectively, then, for Problem (30), \mathcal{L} is given as

$$\mathcal{L}(\theta, \lambda, \mu) = \mathbf{1}'\theta^s - s\psi'\theta - \lambda'\theta + \mu(\mathbf{1}'\theta^s - 1) \quad (31)$$

Setting the Lagrangian's gradient with respect to θ to $\mathbf{0}$ yields

$$\theta = \frac{1}{(\mu + 1)^r} \left(\psi + \frac{1}{s}\lambda \right)^r \quad (32)$$

where $r \triangleq \frac{1}{s-1}$. Now, $\theta \succeq \mathbf{0}$ implies via Equation (32) that $\lambda \succeq -s\psi$. Since $\lambda \succeq \mathbf{0}$ must also hold, we have that $\lambda \succeq \max\{\mathbf{0}, -s\psi\} = \mathbf{0}$, as $\psi \succeq \mathbf{0}$ by assumption. Given this range of λ , \mathcal{L} is maximized for $\lambda^* = \mathbf{0}$. Thus, Equation (32) becomes

$$\theta = \frac{1}{(\mu + 1)^r} \psi^r \quad (33)$$

In view of Equation (33), the last constraint, $\|\theta\|_s^s \leq 1$, implies that $\mu \geq \|\psi\|_{rs} - 1$. Given that $\mu \geq 0$ must also hold, we conclude that $\mu \geq \max\{0, \|\psi\|_{rs} - 1\}$. Finally, for this range of μ , \mathcal{L} is maximized for $\mu^* = \max\{0, \|\psi\|_{rs} - 1\}$ and, therefore, substituting μ^* into Equation (33) finally yields the desired minimizer Equation (15). Note that the results we found for θ^* , λ^* and μ^* satisfy the problem's Karush-Kuhn-Tucker (KKT) conditions.

Next, we prove the second part of the theorem, where $\Omega(\theta) \triangleq \{\theta \mid \theta \succeq \mathbf{0}, \theta \in \bar{B}_{\theta,1}\}$. In this case, the Lagrangian is given as

$$\mathcal{L}(\theta, \lambda, \mu) = \frac{1}{s}\mathbf{1}'\theta^s - \psi'\theta - \lambda'\theta + \mu(\mathbf{1}'\theta - 1) \quad (34)$$

By setting the gradient with respect to θ to $\mathbf{0}$ yields

$$\theta = (\psi + \lambda - \mu\mathbf{1})^r \quad (35)$$

Since $\theta \succeq \mathbf{0}$, and the complementary slackness condition $\lambda_m \theta_m = 0$ must hold, it gives $\lambda = \max\{\mathbf{0}, \mu\mathbf{1} - \psi\}$. Plug it into Equation (34), we have

$$\begin{aligned} \theta &= (\psi + \max\{\mathbf{0}, \mu\mathbf{1} - \psi\} - \mu\mathbf{1})^r \\ &= (\max\{\psi - \mu\mathbf{1}, \mathbf{0}\})^r \end{aligned} \quad (36)$$

To maximize the Lagrangian with respect to μ , the latter variable should be as small as possible. When combining the conditions $\mu \geq 0$ and $\|\theta\|_1 \leq 1$, we reach the desired conclusion.

.4 Proof to Lemma 3

First note that, when $p \in (0, \frac{1}{2})$, we have $\frac{1}{q} \in (1, +\infty)$. We first form the Lagrangian of Problem (20):

$$\mathcal{L}(\phi, \alpha) = \mathbf{g}'\phi^{-\frac{1}{q}} + \alpha(\sum_{t=1}^T \phi^t - 1) - \sum_{t=1}^T \beta^t \phi^t \quad (37)$$

Setting the derivative of Equation (37) with respect to ϕ^t equal to 0 yields:

$$\phi^t = \left[\frac{g^t}{(\alpha - \beta^t)q} \right]^p, \quad \forall t \quad (38)$$

Obviously, as long as $g^t \neq 0$, we have $\phi^t \neq 0$ and therefore $\beta^t = 0$ due to the complementary condition $\beta^t \phi^t = 0, \forall t$. In this situation, we have that $\phi^t \propto (g^t)^p$. On the other hand, if $g^t = 0$, there must be $\phi^t = 0$. Therefore, in summary, we have $\phi \propto \mathbf{g}^p$. Since we can observe that the optimum ϕ^* must lie on the boundary, where $\sum_{t=1}^T \phi^t = 1$, we have

$$\phi^* = \frac{\mathbf{g}^p}{\sum_{t=1}^T (g^t)^p} \quad (39)$$

which completes the proof.